IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellants: Robert Jochemsen et al. Group Art Unit: 2186

Serial No.: 10/533,735 Examiner: Schnee, Hal W.

Filed: May 4, 2005 Confirmation No.: 4556

For: METHOD AND DEVICE FOR PERSISTENT-MEMORY MANAGEMENT

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF UNDER 37 C.F.R. § 41.37(a)

This is an appeal to the Board of Patent Appeals and Interferences from the decision of the Examiner dated August 7, 2008, which finally rejected claims 1-8 and 10-25 in the above-identified application. The Office date of receipt of Appellants' Notice of Appeal was November 7, 2008. This Appeal Brief is hereby submitted pursuant to 37 C.F.R. § 41.37(a).

CERTIFICATE OF MAILING UNDER 37 C.F.R. 1.8
I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being facsimile transmitted to the Patent and Trademark Office on the date shown below.

Date:_____ Signed:_____
Typed Name: Mark A. Wilson

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST

The real party in interest is the assignee of the full interest in the invention, NXP B.V., of Eindhoven, Netherlands.

II. RELATED APPEALS AND INTERFERENCES

To the best of Appellants' knowledge, there are no appeals or interferences related to the present appeal that will directly affect, be directly affected by, or have a bearing on the Board's decision in the instant appeal.

III. STATUS OF CLAIMS

Claims 1-8 and 10-25 are pending in the present application.

Claim 9 is canceled.

No claims are withdrawn.

No claims are objected to.

Claims 1-8 and 10-25 stand rejected as follows:

Claims 1-4, 7, 8, and 10 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale (U.S. Pat. No. 6,804,763, hereinafter Stockdale), in view of Cheng (U.S. Pat. No. 5,701,516, hereinafter Cheng), and further in view of O'Neill (U.S. Pat. Pub. No. 2003/0182414, hereinafter O'Neill);

Claims 5, 6, and 11-15 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale, in view of Cheng, further in view of O'Neill, and further in view of Hanes (U.S. Pat. Pub. No. 2003/0081932, hereinafter Hanes); and

Claims 16-25 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Stockdale in view of Cheng in view of O'Neill in view of Hanes and further in view of Lee (U.S. Pat. No. 5,930,167, hereinafter Lee).

Claims 1-8 and 10-25 are the subject of this appeal. A copy of claims 1-8 and 10-25 is set forth in the Claims Appendix.

IV. STATUS OF AMENDMENTS

Proposed amendments were submitted subsequent to the Final Office Action mailed August 7, 2008. The proposed amendments were submitted to place claims 1-8 and 10-25 in better condition for appeal. However, the proposed amendments were not entered by the Examiner because the proposed amendments purportedly raise new issues that would require a further consideration and/or search, according to the Advisory Action.

V. SUMMARY OF CLAIMED SUBJECT MATTER

This section of this Appeal Brief is set forth to comply with the requirements of 37 C.F.R. § 41.37(c)(1)(v) and is not intended to limit the scope of the claims in any way. Examples of implementations of the limitations of independent claims 1, 5, 7, 10, and 11 are described below.

The language of claim 1 relates to a memory management device for managing a memory space of at least one persistent-memory device. Page 1, lines 1-5. In particular, claim 1 recites a memory allocation unit to communicate with at least one application device and to allocate at least one first part of the memory space to the application device to write a first working data structure with working data blocks to the memory space and to write a second working data structure as a copy of the working data blocks. Page 12, lines 21-28. The second working data structure is a copy of the first working data structure in the same memory space as the first working data structure. Page 16, lines 19-26. The allocation unit communicates with at least one file system device, and allocates the first part of the memory space on request from the application device or from the file system device. Page 12, lines 21-28. The persistent-memory device is used as a write cache memory for the file system device. Page 4, lines 14-16.

The language of claim 5 relates to a file system device with a file allocation unit that maintains a file allocation table at a current status. Page 13, lines 11-17. The file allocation table assigns at least one disk space address to at least one file. Page 13, lines 11-17. The file allocation unit communicates with a memory management device that is related to a persistent-memory device and includes an address of at least one first memory space of the persistent-memory device in the maintenance of the file allocation

table. Page 13, lines 11-17. The file system device converts a copy of the first working data structure to a file data structure and writes the file data structure to a secondary storage medium. Page 8, lines 5-15. The copy of the first working data structure is written to the same persistent-memory device as the first working data structure. Page 8, lines 5-15. The persistent-memory device is used as a write cache memory for the file system device. Page 19, lines 20-25.

The language of claim 7 relates to an application device with a persistent-memory device connected to a processor and a data management unit that manipulates data in the persistent-memory device. Page 13, lines 27-34. The data management unit writes at least one third executable file to the persistent-memory device, or to provide the file system with a reference to at least one third executable file in the file system. Page 13, lines 27-34. By executing the third executable file the processor transforms a copy of a first working data structure into a predetermined data-sequence form and the copy of the first working data structure is stored in a same persistent-memory device as the first working data structure. Page 8, lines 5-15. The persistent-memory device is used as a write cache memory for a file system device. Page 19, lines 20-25.

The language of claim 10 relates to a data processing system with a memory management device for managing a memory space of at least one persistent-memory device. Page 12, lines 21-28. The memory management device has a memory allocation unit that communicates with at least one application device and to allocate at least one first part of the memory space to the application device to write a first working data structure of working data blocks to the memory space. Page 12, lines 21-28. A second working data structure is a copy of the working data blocks, and the second working data structure is a copy of the first working data structure in the same memory space as the first working data structure. Page 12, lines 21-28. The allocation unit communicates with at least one file system device and allocates on request from the application device or from the file system device the first part of the memory space to the file system. Page 12, lines 21-28. The persistent-memory device is used as a write cache memory for the file system device. Page 19, lines 20-25.

The language of claim 11 relates to a method for managing memory space of a persistent-memory device, including allocating at least one first part of the memory space

to a file system device and writing a first working data structure of working data blocks to the memory space. Page 12, lines 21-28. The second working data structure is a copy of the working data blocks, and the copy is in the same memory space as the first working data structure. Page 8, lines 5-15. The second working data structure is converted into a predetermined data-sequence form. Page 13, lines 27-34. The persistent-memory device is used as a write cache memory for the file system device. Page 19, lines 20-25.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Whether claims 1-4, 7, 8, and 10 are patentable over Stockdale, Cheng, and O'Neill under 35 U.S.C. 103(a).

B. Whether claims 5, 6, and 11-15 are patentable over Stockdale, Cheng, O'Neill, and Hanes under 35 U.S.C. 103(a).

C. Whether claims 16-25 are patentable over Stockdale, Cheng, O'Neill, Hanes, and Lee under 35 U.S.C. 103(a).

VII. ARGUMENT

For the purpose of this appeal, claims 1-4, 7, 8, and 10 are argued together as a group for the purposes of the question of patentability over the combination of Stockdale, Cheng, and O'Neill under 35 U.S.C. 103(a). Claims 5, 6, and 11-15 are argued as a separate group for purposes of the question of patentability over the combination of Stockdale, Cheng, O'Neill, and Hanes under 35 U.S.C. 103(a). Claims 16-25 are argued as a separate group for purposes of the question of patentability over the combination of Stockdale, Cheng, O'Neill, Hanes, and Lee under 35 U.S.C. 103(a).

A. Claims 1-4, 7, 8, and 10 are patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of references does not teach all the limitations of the claims.

Appellants respectfully submit that claim 1 is patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of references does not teach all of the limitations of the claim. Claim 1 recites:

A memory management device for managing a memory space of at least one persistent-memory device, comprising:

a memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said persistent-memory device to said application device to write a first working data structure comprising a plurality of working data blocks to the persistent-memory device and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same persistent-memory device as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said persistent-memory device to said file system; wherein the persistent-memory device is used as a write cache memory for said file system device

(Emphasis added).

In contrast, the combination of Stockdale, Cheng, and O'Neill does not teach all of the limitations recited in the claim because the combination of Stockdale, Cheng, and O'Neill does not teach writing a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same persistent-memory device as the first working data structure, as recited in the claim. Additionally, the combination of Stockdale, Cheng, and O'Neill does not teach an allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said persistent-memory device to said file system.

The Examiner relies on O'Neill as purportedly teaching writing a second working data structure in the same memory space as the first working data structure. Also, the Examiner relies on Stockdale as purportedly teaching an allocation unit that is adapted to communicate with at least one file system. In particular, O'Neill and Stockdale fail to teach the indicated limitations because 1) O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure, and 2) Stockdale does not teach an allocation unit that is adapted to communicate with at least one file system.

1. O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure.

O'Neill was cited by the Office Action dated 8/7/08 as teaching writing a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure. However, O'Neill does not teach a copy of the working data blocks in the same memory space as the first working data structure. O'Neill identifies the working data bank as being volatile memory and the copy or backup bank as being non-volatile memory. O'Neill, page 17, paragraph 133.

As explained in previous responses, O'Neill does not teach having a copy of the working data blocks in the same memory space as the first working data structure. As described by O'Neill, the bank 1234 of the flash memory 1002 may be configured to store a copy of the contents of the bank 1232 of the RAM 1004. In other words, as described in O'Neill, each of the banks 1232 and 1234 are separate allocations from one another. Additionally, the banks 1232 and 1234 are allocated in different memory devices, in the RAM 1004 and the flash memory1002, respectively. Therefore, none of the data that may be stored in the banks 1232 or 1234 could be stored in the same memory space as each other because the data stored in each bank is stored on a different memory device. Moreover, O'Neill states that the data stored in the bank 1234 is a copy of the data stored in the bank 1232. However, as previously stated, the bank 1232 is allocated in the RAM 1004 and the bank 1234 is allocated in the flash memory 1002. Therefore, not only is the data stored in different allocations with regard to banks 1232 and 1234, but the data is stored in different memory devices. Hence, in regard to the banks 1232 and 1234, O'Neill does not teach writing a first working data structure and a second working data structure to the same memory space. O'Neill merely teaches having a copy of a copy of the working data in a non-volatile memory bank, and does not teach having a working data and a copy of the working data both in the same memory space as cited in the claim.

2. Stockdale does not teach an allocation unit that is adapted to communicate with at least one file system.

Stockdale merely teaches having the data that is stored in non-volatile memory by the application device appear as directories and files. Stockdale, col. 34, lines 51-67. If an allocation unit were to communicate with a file system device it would require the file system to be a separate device from the application device. In other words, there would be no need for an application device to communicate with itself. Therefore, the file system device must be a different device from the application device. Stockdale teaches that the application device formats the critical data that it stores in non-volatile memory as file structures to make it easy for standard operating systems and application tools to access the data. Stockdale, col. 35, lines 1-5. This is merely a first working data structure with working data blocks stored in non-volatile memory by the application device, not data stored on a file system device.

In addition, if the allocation unit as described by Stockdale is to communicate with at least one file system device, the file system device must have a device driver for the file system device. A review of the device driver list in FIG. 2 does not show a device driver for a file system device. Also, the gaming machine system shown in FIG. 2 in Stockdale does not show a file system device in the list of device interfaces. Stockdale merely teaches storing files of critical game data in non-volatile memory rather than communicating with a file system device and allocating memory on request from said file system device as cited in the claim. Stockdale teaches an event manager that controls access to the NV-RAM where the critical data is stored. Stockdale, Col. 10, lines 46-67. Stockdale also teaches that various devices can request access to the NV-RAM. Stockdale, Col.11, lines 1-12. But Stockdale does not teach that the event manager can or does interface with a file system device.

Therefore, the combination of Stockdale, Cheng, and O'Neill does not teach the limitations of claim 1 because the combination of Stockdale, Cheng, and O'Neill does not teach writing a second working data structure in the same memory space as the first working data structure. Additionally, the combination of Stockdale, Cheng, and O'Neill does not teach an allocation unit that is adapted to communicate with at least one file

system. Accordingly, Appellants respectfully assert claim 1 is patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of Stockdale, Cheng, and O'Neill does not teach all of the limitations of the claim. Consequently, Appellants request that the rejection of claim 1 under 35 U.S.C. 103(a) be withdrawn.

Appellants respectfully assert independent claim 7 is also patentable over Stockdale, Cheng, and O'Neill at least for similar reasons to those stated above in regard to the rejection of independent claim 1. In particular, the rejection of claim 7 relies on similar reasoning to characterize the writing of a second working data structure in the same memory space as the first working data structure. Here, although the language of claim 7 differs from the language of claim 1, and the scope of claim 7 should be interpreted independently of claim1, Appellants respectfully assert that the remarks provided above in regard to claim 1 also apply to the rejection of claim 7. Accordingly, Appellants respectfully assert independent claim7 is also patentable over the combination of Stockdale, Cheng, and O'Neill because the combination of Stockdale, Cheng, O'Neill, and Hanes does not teach all the limitations of the claim. Consequently, Appellants request that the rejection of claim 7 under 35 U.S.C. 103(a) be withdrawn.

Appellants respectfully assert independent claim 10 is also patentable over Stockdale, Cheng, O'Neill, and Hanes at least for similar reasons to those stated above in regard to the rejection of independent claim 1. In particular, the rejection of claim 10 relies on similar reasoning to characterize the writing of a second working data structure in the same memory space as the first working data structure. Here, although the language of claim 10 differs from the language of claim 1, and the scope of claim10 should be interpreted independently of claim1, Appellants respectfully assert that the remarks provided above in regard to claim 1 also apply to the rejection of claim 10. Accordingly, Appellants respectfully assert independent claim10 is also patentable over the combination of Stockdale, Cheng, O'Neill, and Hanes, because the combination of Stockdale, Cheng, O'Neill, and Hanes does not teach all the limitations of the claim. Consequently, Appellants request that the rejection of claim 10 under 35 U.S.C. 103(a) be withdrawn.

Given that claims 2-4 and 8 depend from and incorporate all of the limitations of the independent claims 1 and 7, which are patentable over the cited references,

Appellants respectfully submit that dependent claims 2-4 and 8 are also patentable over the cited reference based on allowable base claims. Additionally, each of claims 2-4 and 8 may be allowable for further reasons. Consequently, Appellants request that the rejections of claims 2-4 and 8 under 35 U.S.C. 103(a) be withdrawn.

B. Claims 5, 6, and 11-15 are patentable over the combination of Stockdale, Cheng, O'Neill, and Hanes because the combination of references does not teach all the limitations of the claims.

Appellants respectfully submit that claim 5 is patentable over the combination of Stockdale, Cheng, O'Neill, and Hanes because the combination of references does not teach all of the limitations of the claim. Claim 5 recites:

> A file system device comprising a file allocation unit adapted to maintain a file allocation table at a current status, said file allocation table assigning at least one disk space address to at least one file, wherein said file allocation unit is adapted to communicate with a memory management device that is related to a persistent-memory device and to include an address of at least one first memory space of said persistent-memory device in the maintenance of said file allocation table, wherein the file system device is configured to convert a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium, wherein the copy of the first working data structure is written to a same memory space as the first working data structure, and wherein the persistent memory is used as a write cache memory for said file system device

(Emphasis added).

In contrast, the combination of Stockdale, Cheng, O'Neill, and Hanes does not teach all of the limitations recited in the claim because the combination of Stockdale, Cheng, O'Neill, and Hanes does not teach a file allocation table assigning at least one disk space address to at least one file. Also, the combination of references does not teach converting a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium. Also, the combination of references does not teach a copy of the first working data structure is written to a same memory space as the first working data structure.

The Examiner relies on O'Neill as purportedly teaching writing a second working data structure in the same memory space as the first working data structure, as recited in the claim. Also, the Examiner relies on Stockdale as purportedly teaching an allocation

unit that is adapted to communicate with at least one file system. In particular, O'Neill and Stockdale fail to teach the limitation of the claims, because 1) O'Neill does not teach converting a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium, 2) Stockdale does not teach assigning at least one disk space address to at least one file, and 3) Stockdale does not teach the copy of the first working data structure is written to a same memory space as the first working data structure.

1. O'Neill does not teach converting a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium.

In the rejection of claim 5, the Office Action relies on O'Neill to teach converting a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium. However, O'Neill only teaches writing a second working data structure to flash memory as a copy of the first working structure. O'Neill is silent on converting the copy of the first data structure to a file data structure and to write the file data structure to a secondary storage medium. O'Neill, page 18, paragraph 142. Moreover, secondary storage media is different than a primary storage media because secondary storage media is not directly accessible to a computer and the data to be written to secondary storage must be converted to a file data structure in order to be written to the secondary storage media. Flash memory is considered primary storage because it is directly addressable by the computer. Therefore, because O'Neill only teaches writing data to flash memory it does not teach the file system device is configured to convert a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium.

2. Stockdale does not teach assigning at least one disk space address to at least one file.

In the rejection of claim 5, Stockdale is cited as purportedly teaching assigning at least one disk space address to at least one file. In contrast to the language of claim 5, Stockdale does not teach a file allocation unit that assigns at least one disk space address to at least one file because Stockdale does not teach any form of disk storage device

associated with the gaming system. Stockdale is completely silent on interfacing with a disk system and allocating memory for a disk system. Rather, Stockdale merely teaches utilizing non-volatile memory to create a non-volatile memory file system. Stockdale, col. 34, lines 51-58. Additionally, Stockdale merely teaches about assigning operating system permissions to the non-volatile memory system files. Stockdale, col. 35, lines 14-20. Neither of these teachings relates to assigning at least one disk space address to at least one file, as cited in the claims of the present application.

3. Stockdale does not teach the copy of the first working data structure is written to a same memory space as the first working data structure.

The argument for this limitation relies on similar reasons to those stated above in regard to the rejection of this limitation of independent claim 1. In particular, the rejection of claim 5 relies on similar reasoning to characterize the copy of the first working data structure is written to a same memory space as the first working data structure. Here, although the language of claim 5 differs from the language of claim 1, and the scope of claim 5 should be interpreted independently of claim 1, Appellants respectfully assert that the remarks provided above in regard to the rejection of claim 1 also apply to the rejection of this limitation of claim 5. Additionally, the rejection of claim 5 does not rely on Cheng, O'Neill, or Hanes as possibly teaching the missing limitations of Stockdale. Consequently, Appellants request that the rejection of claim 5 under 35 U.S.C. 103(a) be withdrawn

Appellants respectfully assert independent claim 11 is also patentable over Stockdale, Cheng, O'Neill, and Hanes at least for similar reasons to those stated above in regard to the rejection of independent claim 5. In particular, the rejection of claim 11 relies on similar reasoning to characterize wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure of O'Neill. Here, although the language of claim 11 differs from the language of claim 5, and the scope of claim 11 should be interpreted independently of claim 5, Appellants respectfully assert that the remarks provided above in regard to claim 5 also apply to the rejection of claim 11. Accordingly, Appellants respectfully assert independent claim11 is also patentable over the combination of Stockdale, Cheng,

O'Neill, and Hanes, because the combination of Stockdale, Cheng, O'Neill, and Hanes does not teach all the limitations of the claim. Consequently, Appellants request that the rejection of claim 11 under 35 U.S.C. 103(a) be withdrawn.

Given that claims 6 and 12-15 depend from and incorporate all of the limitations of the independent claims 5 and 11, which are patentable over the cited references, Appellants respectfully submit that dependent claims 6 and 12-15 are also patentable over the cited reference based on allowable base claims. Additionally, each of claims 6 and 12-15 may be allowable for further reasons. Consequently, Appellants request that the rejections of claims 6 and 12-15 under 35 U.S.C. 103(a) be withdrawn.

C. Claims 16-25 are patentable over Stockdale, Cheng, O'Neill, Hanes, and Lee because the combination of references does not teach all the limitations of the claims.

Given that claims 16-25 depend from and incorporate all of the limitations of the independent claim 11, which is patentable over the cited references, Appellants respectfully submit that dependent claims 16-25 are also patentable over the cited references based on an allowable base claim. Additionally, each of claims 16-25 may be allowable for further reasons. Consequently, Appellants request that the rejections of claims 16-25 under 35 U.S.C. 103(a) be withdrawn

VIII. CONCLUSION

For the reasons stated above, claims 1-8 and 10-25 are patentable over the cited references. Thus, the rejections of claims 1-8 and 10-25 should be withdrawn. Appellants respectfully request that the Board reverse the rejections of claims 1-8 and 10-25 under 35 U.S.C. 103(a) and, since there are no remaining grounds of rejection to be overcome, direct the Examiner to enter a Notice of Allowance for claims 1-8 and 10-25.

At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account **50-3444** pursuant to 37 C.F.R. 1.25. Additionally, please charge any fees to Deposit Account **50-3444** under 37 C.F.R. 1.16, 1.17, 1.19, 1.20 and 1.21.

Respectfully submitted,

/mark a. wilson/

Date: January 7, 2009

Mark A. Wilson
Reg. No. 43,994

Wilson & Ham
PMB: 348
2530 Berryessa Road
San Jose, CA 95132
Phone: (925) 249-1300
Fax: (925) 249-0111

IX. CLAIMS APPENDIX

1. A memory management device for managing a memory space of at least one persistent-memory device, comprising:

a memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said memory space to said application device to write a first working data structure comprising a plurality of working data blocks to the memory space and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said memory space to said file system; wherein the persistent memory is used as a write cache memory for said file system device.

2. A memory management device according to claim 1, wherein said memory allocation unit is adapted to maintain a memory allocation table at a current status, said memory allocation table assigning at least one memory address representing a defined part of said memory space to either said application device or to said file system device.

3. A memory management device according to claim 2, further comprising a processor and a memory, wherein said memory allocation unit is implemented in the form of at least one first executable file contained in said memory.

4. A memory management device according to claim 3, wherein said memory is a persistent-memory device, in particular said persistent-memory device.

5. A file system device comprising a file allocation unit adapted to maintain a file allocation table at a current status, said file allocation table assigning at least one disk space address to at least one file, wherein said file allocation unit is adapted to communicate with a memory management device that is related to a persistent-memory

device and to include an address of at least one first memory space of said persistent-memory device in the maintenance of said file allocation table, wherein the file system device is configured to convert a copy of a first working data structure to a file data structure and to write the file data structure to a secondary storage medium, wherein the copy of the first working data structure is written to a same memory space as the first working data structure, and wherein the persistent memory is used as a write cache memory for said file system device.

6. A file system device according to claim 5, further comprising a processor and a memory, wherein said file allocation unit is implemented in the form of at least one second executable file contained in said memory.

7. An application device, comprising a persistent-memory device connected to a processor, and a data management unit adapted to manipulate data in said persistent memory device, wherein said data management unit is adapted to write at least one third executable file to said persistent memory device, or to provide the file system with a reference to at least one third executable file in said file system, such that by executing said third executable file said processor is adapted to transform a copy of a first working data structure into a predetermined data-sequence form, wherein the copy of the first working data structure is stored in a same memory space as the first working data structure; and wherein the persistent memory is used as a write cache memory for a file system device.

8. An application device according to claim 7, wherein said data management unit is provided in the form of least one fourth executable file in a memory, particularly, in said persistent memory.

9. (canceled)

10. A data processing system, comprising a memory management device for managing a memory space of at least one persistent-memory device, comprising a

memory allocation unit adapted to communicate with at least one application device and to allocate at least one first part of said memory space to said application device to write a first working data structure comprising a plurality of working data blocks to the memory space and to write a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure, wherein said allocation unit is further adapted to communicate with at least one file system device, and to allocate on request from said application device or from said file system device said first part of said memory space to said file system; wherein the persistent memory is used as a write cache memory for said file system device.

11. A method for managing memory space of a persistent-memory device, comprising:

allocating at least one first part of said memory space to a file system device;

writing a first working data structure comprising a plurality of working data blocks to the memory space;

writing a second working data structure comprising a copy of the plurality of working data blocks, wherein the second working data structure comprises a copy of the first working data structure in the same memory space as the first working data structure; and

converting the second working data structure into a predetermined data-sequence form;

wherein the persistent memory is used as a write cache memory for said file system device.

12. A method according to claim 11, wherein said allocating step comprises a step of blocking a writing access to said first part of said memory space.

13. A method according to claim 12, wherein said allocating step comprises a step of giving away to said file system device the power of reading access to said first part of said memory space.

14. A method according to claim 11, comprising a step of deallocating said first part of said memory space to a memory management device.

15. A method according to claim 14, wherein said allocating step or said deallocating step comprises transmitting an address range defining said first part of said memory space from said memory management device to said file system device or, respectively, vice versa.

16. A method according to claim 14, wherein said deallocating step is performed for said first part of said memory space given the condition that first data contained in said first part of said memory space is stored in the form of file data in a second part of said memory space, said file data having a predetermined file structure, and that said second part of said memory space is allocated to said file system device.

17. A method according to claim 16, wherein said deallocating step is performed for said second part of said memory space given the condition that said file data has been written to a secondary storage medium.

18. A method for write-caching first data worked on by an application, said first data being contained in a first part of a memory space of a persistent-memory device, comprising a step of performing a memory managing method according to claim 17.

19. A write-caching method according to claim 18, comprising, after said allocating step, a step of sending a confirmation message from said file system device to said application device.

20. A write-caching method according to claim 18, wherein said first data is a copy of third data contained in a third part of said memory space, said write-caching method comprising, before performing said memory managing method, a step of copying said third data to said first memory space.

21. A write-caching method according to claim 18, comprising the steps of

allocating a fourth part of said memory space to said application device for an executable file or dynamic link library that is adapted to converting said first data into file data

writing said executable file or dynamic link library to said fourth part of said memory space

allocating said fourth part of said memory space to said file system device.

22. A write-caching method according to claim 21, comprising a step of transforming said first data into said file data with the aid of said executable file or said dynamic link library.

23. A write-caching method according to claim 22, wherein said transforming step is initiated by said file system device.

24. A write-caching method according to claim 23, comprising, after said transforming step, a step of deallocating said fourth part of said memory space to said memory management device.

25. A method for saving data worked on by an application device to a file on a secondary storage medium, comprising performing a write-caching method according to claim 18, and further comprising writing said file data to said secondary storage medium.

X. EVIDENCE APPENDIX

There is no evidence submitted with this Appeal Brief.

XI. RELATED PROCEEDINGS APPENDIX

To the best of Appellants' knowledge, there are no appeals or interferences related to the present appeal that will directly affect, be directly affected by, or have a bearing on the Board's decision in the instant appeal.